Motion Estimation via Robust Decomposition with Constrained Rank

German Ros*, Jose M. Álvarez, and Julio Guerrero

arXiv:1410.6126v1 [cs.CV] 22 Oct 2014

***Abstract*—In this work, we address the problem of outlier detection for robust motion estimation by using modern sparse-low-rank decompositions, i.e., Robust PCA-like methods, to impose global rank constraints. Robust decompositions have shown to be good at splitting a corrupted matrix into an uncorrupted low-rank matrix and a sparse matrix, containing outliers. However, this process only works when matrices have relatively low rank with respect to their ambient space, a property not met in motion estimation problems. As a solution, we propose to exploit the partial information present in the decomposition to decide which matches are outliers. We provide evidences showing that even when it is not possible to recover an uncorrupted low-rank matrix, the resulting information can be exploited for outlier detection. To this end we propose the Robust Decomposition with Constrained Rank (RD-CR), a proximal gradient based method that enforces the rank constraints inherent to motion estimation. We also present a general framework to perform robust estimation for stereo Visual Odometry, based on our RD-CR and a simple but effective compressed optimization method that achieves high performance. Our evaluation on synthetic data and on the KITTI dataset demonstrates the applicability of our approach in complex scenarios and it yields state-of-the-art performance.**

***Index Terms*—Robust statistics, Robust PCA, Visual Odometry**

I. Introduction

THIS paper addresses the application of robust decomposition techniques, like the successful Robust PCA (RPCA) [1] [2], to the problem of outlier (i.e., gross noise) detection in the special case of camera-pose recovery and Visual Odometry, referred in this paper to as motion estimation.

The emergence of outliers in real situations is unavoidable and estimation techniques need to achieve a certain level of robustness to prevent producing wrong estimations. In the context of motion estimation —the practical case studied in this paper— outliers arise by two main reasons. The first reason is the violation of model assumptions, for instance, when a scene is considered static but foreground objects start moving (see the car in Fig. 1). The second is the failure of matching methods to retrieve correct point matches at different viewpoints.

Recently, it has been shown that RPCA can be used to recover an uncorrupted version of a matrix affected by outliers [1] [2] [3], but such a recovery is just achievable when matrices fulfil several well-understood constraints [4] (see Section III). However, this is not the case when dealing with the data matrices arising in motion estimation problems, and therefore RPCA can not be used to exactly recover an uncorrupted version of the data, i.e., such matrices are outside the recovery boundaries of RPCA [1] [2].

In this paper we propose to exploit the information resultant from the sparse and low-rank decomposition to make a binary decision on each tuple of point matches (column) about its pertinence to the outlier set. As we will show along this paper, despite the impossibility of performing an exact recovery of every element of the observation matrix, the resultant information is enough to make this set of binary decisions. Unlike works like [3], our outlier detection technique has the advantage that is not based on the correct recovery of the low-rank matrix and the sparse matrix (or their column spaces). Instead, we exploit the rank constraints present in motion estimation problems, leading to a novel Robust Decomposition with Constrained Rank (RD-CR). Our technique is based on proximal gradients, inspired by the Accelerated Proximal Gradient algorithm (APG) [5], but extended to deal with fixed-rank constraints efficiently. We show that our approach presents extended boundaries of application, making robust decomposition techniques useful for motion estimation applications.

We also present a new framework for robust motion estimation that applies RD-CR as a preprocessing stage, prior to the application of a very fast compressed least-squares method. Compressed least-squares makes use of an efficient data compression scheme known as the Reduced Measurement Matrix (RMM) [6] [7] to condense thousands of constraints (point matches) into a compact matrix surrogate that allows for very fast optimization. RD-CR mitigates the robustness problem of compressed least-squares methods by filtering out the outliers. This finally leads to an efficient and robust regression framework. We demonstrate this framework for the stereo Visual Odometry problem with evaluations in both synthetic cases and the realistic and challenging sequences of the KITTI dataset [8]. This evaluation also serves to test the applicability of RD-CR technique in complex scenarios, showing its good virtues when compared against state-of-the-art techniques. To the best of our knowledge this is the first time robust decompositions are tested in such challenging scenarios.

II. Related Work

A. Robust Estimation

Robust estimation aims at providing algorithms with certain level of resistance to noise and outliers [9] [10]. It is a well-studied problem within the field of statistics and plays a

G. Ros is with the Computer Vision Center at Universitat Autònoma de Barcelona, Spain, e-mail: gros@cvc.uab.es .

Jose M. Álvarez is a researcher at NICTA, 2601 Canberra, Australia, e-mail: Jose.Alvarez@nicta.com.au.

J. Guerrero is with the Department of Applied Mathematics at Universidad de Murcia, Spain, e-mail: juguerre@um.es .

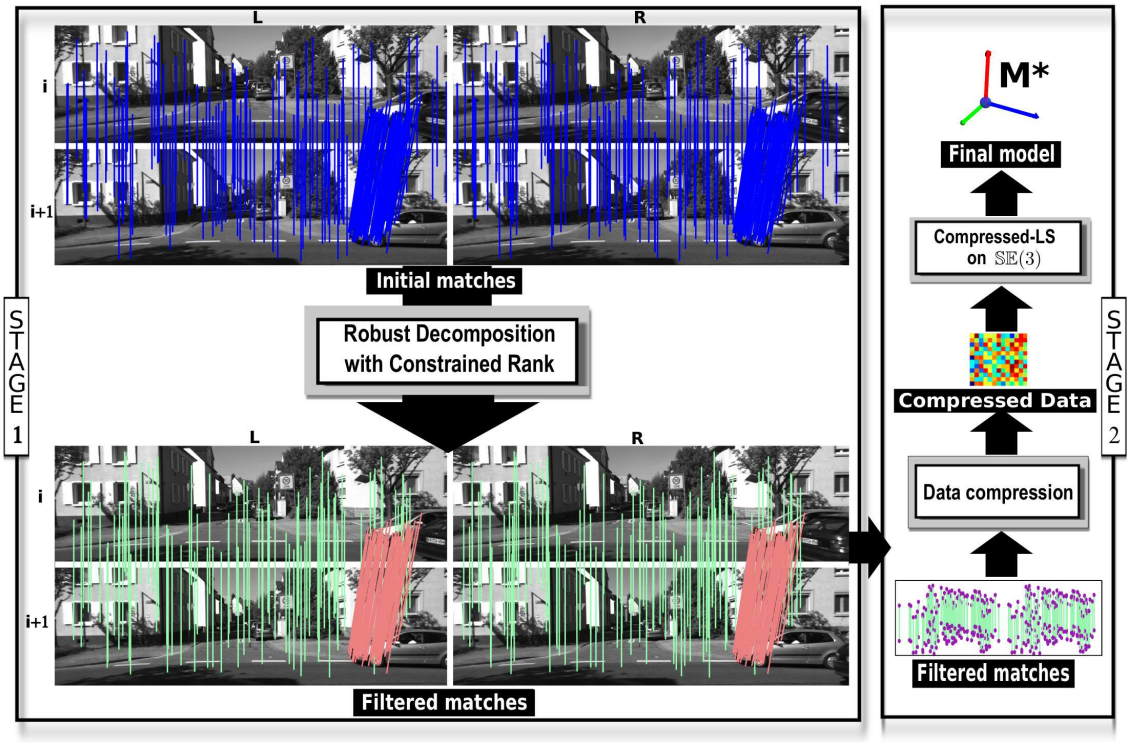

Fig. 1. Proposed approach as a two stages pipeline, i.e.: RD-CR for outlier detection followed by a stage of data reduction and optimization. Notice how the method can correctly identify the matches from the moving car as outliers during stage one.

crucial role in computer vision and robotics. Within computer vision and robotics, these techniques are widely applied to motion estimation problems, such as camera pose estimation and Visual Odometry (VO). In these domains, state-of-the-art approaches are usually based on consensus strategies, like the RANdom SAmple Consensus (RANSAC) [11] or any of its variants [9]. The main problem of RANSAC is its computational complexity. RANSAC estimates robust models by identifying inliers via the generation of several hypotheses (models). This process grows exponentially w.r.t the percentage of outliers, leading to undesired performance for practical applications. A recent variation of consensus methods, called robust averaging [12] [13], avoids this drawback by merging several "weak" models in a robust fashion to produce a robust outcome. In this case less models are needed, leading to a more efficient approach. However these methods have a breakdown point of $\beta = 0.5$.

In many occasions, simpler alternatives, such as robust cost functions, are required to perform a quick and simple job. An instance of this is the Huber M-estimator [14] [10], a variation of least-squares that uses two different modes to weight data points close to zero (inliers) and those that are far away (outliers) . Unfortunately, although efficient, its theoretical breakdown point is $\beta = 0$, meaning that in the worst case a single gross error could produce an arbitrarily large bias in the estimation. Another example is the Iteratively Reweighted Least Squares (IRLS) technique [15] [16]. In this case robustness is achieved by dynamically modifying a set of weights describing the influence of each observation. When this update is based on a robust cost like the ℓ_1-norm (IRLS-L1) the breakdown point becomes $\beta = 0.5$, therefore, being a good alternative to consensus approaches [15].

Achieving a breakdown point of $\beta = 0.5$ is the best one can expect without any extra information of the data distribution. In previous cases, all the methods try to detect outliers by analysing how data samples agree with a set of putative models. This exhaustive generation leads to an expensive process. A promising alternative we use in this paper consists in analysing the quality of the data with respect to an abstract set of constraints, such as the rank constraints arising in $\mathbb{SO}(N)$, $\mathbb{SE}(N)$ groups. These rank constraints can be exploited very efficiently without the need of creating a set of putative models. Instead, observations are checked as a "whole" against the rank constraint to decide which elements are more likely to violate the condition. Moreover, these properties still hold beyond the breakdown point $\beta = 0.5$, at least theoretically.

B. RPCA for Outlier Detection

Outlier detection has been previously addressed by RPCA-like techniques in different domains, such as identifying outliers in digits images [3]. However, in these cases the ambient dimension is high and the rank of the data matrix is low (i.e., the recovery warranties are met). As a consequence, outlier detection is posed as the problem of recovering a clean column space of the input data by mean of a robust decomposition based on a $\ell_{1,2}$ penalizer. However, the problem of motion estimation from stereo devices involves dealing with a low ambient dimension problem, which is very challenging due to theoretical limitations.

An initial attempt to deal with this problem can be found in [17]. The proposed algorithm performs outlier detection for

3D registration by using a RPCA-like approach. However, they always stay within the recovery boundaries by gathering as many frames as possible, creating a high ambient dimension. This avoids the main problem of the low ambient dimension but at the cost of limiting the applicability of the technique, which can only work in small scenarios where conditions remain similar through time.

In this paper we propose to overcome these limitations with a more general alternative that is suitable for large sequences in outdoor environments, where only pairs of frames are required. The recovery boundaries are not fulfilled, but even in these extreme situations, our approach produces useful information that helps detecting outliers. Moreover, we assume that our models are based on multiple 2D observations [18] [19] instead of 3D. This has proven to be more convenient for Visual Odometry, where, least-squares based methods are frequently used to model the problem and the Maximum Likelihood Estimator (MLE) of these functions is achieved by using 2D observations.

III. Problem formulation

In this work, we consider the estimation of motion models $\{M_i\}_{i=1}^{N-1}$, between pair of stereo frames $F_i \rightarrow F_{i+1}$ along a given sequence of N frames $\{F_i\}_{i=1}^{N}$. Each frame $F_i = (V_i^l, V_i^r)$ consists of two images taken from the left and right cameras at the same time instant t_i. This formulation is suitable for the stereo Visual Odometry problem addressed in this work. To this end, we gather observations from two cameras in two different time instants. Camera models are represented by C_i^l, $C_i^r, C_{i+1}^l, C_{i+1}^r$, where superindices stand for (l)eft and (r)ight cameras and subindices denote time instants. Without lack of generality, we assume that the stereo rig is calibrated and that the image planes of both cameras, C_*^l and C_*^r, are aligned and have a known baseline B. The motion of the stereo rig from time t_i to t_{i+1} is given by a rigid 3D transformation M_i as shown in (1),

$$(C_i^l, C_i^r) \overset{M_i \in \mathbb{SE}(3)}{\longrightarrow} (C_{i+1}^l, C_{i+1}^r). \tag{1}$$

This motion is represented as an element of the Special Euclidean group in 3D as in

$$M_i \in \mathbb{SE}(3) = \begin{bmatrix} R_{3\times3} & T_{3\times1} \\ 0_{1\times3} & 1 \end{bmatrix}, \tag{2}$$

where R is a rotation matrix and T a translation vector. Each of the cameras can be algebraically described as shown in (3)–(6),

$$C_i^l = \Pi K \begin{bmatrix} I_{3\times3} & 0_{3\times1} \\ 0_{1\times3} & 1 \end{bmatrix} \tag{3}$$

$$C_{i+1}^l = \Pi K \begin{bmatrix} I_{3\times3} & 0_{3\times1} \\ 0_{1\times3} & 1 \end{bmatrix} M_i \tag{4}$$

$$C_i^r = \Pi K \begin{bmatrix} I_{3\times3} & -B_{3\times1} \\ 0_{1\times3} & 1 \end{bmatrix} \tag{5}$$

$$C_{i+1}^r = \Pi K \begin{bmatrix} I_{3\times3} & -B_{3\times1} \\ 0_{1\times3} & 1 \end{bmatrix} M_i, \tag{6}$$

given that $\Pi(\cdot) : \mathbb{R}^4 \rightarrow \mathbb{R}^3$, represents a projection matrix, and

$$K = \begin{bmatrix} f & 0 & c_u & 0 \\ 0 & f & c_v & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{7}$$

is a calibration matrix with focal length f and principal point at (c_u, c_v). Given this configuration, the transformation M_i of the stereo rig can be estimated from a set of N_c 2D correspondences $\{(x_i^l \leftrightarrow x_i^r \leftrightarrow x_{i+1}^l \leftrightarrow x_{i+1}^r)^{(j)}\}_{j=1}^{N_c}$ coming from the views $V_i^l, V_i^r, V_{i+1}^l, V_{i+1}^r$, respectively [18] [20] [21]. When estimating the transformation M_i, one should account for the presence of noise and outliers in the observations in order to avoid a biased solution.

Our proposal exploits the rank constraints present in rigid 3D motions to identify outliers. To achieve this, it is important to notice that some arrangements of the input data as a matrix lead to such rank constraints [17] [22]. As an example of this concept, let us consider two sets of homogeneous 3D points at two time instants $\mathcal{X}_i, \mathcal{X}_{i+1}$, stacked as described in (8)-left. When the right side expresses a matrix in terms of an initial set of points and two motion transformations acting on it, as

$$\underbrace{\begin{bmatrix} X_i^1 & \cdots & X_i^{N_c} \\ X_{i+1}^1 & \cdots & X_{i+1}^{N_c} \end{bmatrix}}_{\substack{8\times N_c \\ \text{rank}\leq 4}} = \underbrace{\begin{bmatrix} M_{i_{3\times4}} \\ M_{j_{3\times4}} \end{bmatrix}}_{\substack{8\times4 \\ \text{rank}\leq 4}} \underbrace{\begin{bmatrix} X_i^1 & \cdots & X_i^{N_c} \end{bmatrix}}_{\substack{4\times N_c \\ \text{rank}\leq 4}}. \tag{8}$$

From this expression is easy to see that the rank of the matrix is $r \leq 4$, since $\text{rank}(AB) \leq min(\text{rank}(A), \text{rank}(B))$. This example shows the rank constraints exploited by our approach. However, in our case we use multi-view 2D observations from a calibrated stereo rig, instead of 3D points. The reasoning for this case, is given in Lemma 1 and Fig. 2.

Lemma 1. *Let us assume a measurement matrix W made of uncorrupted 2D observations that have been normalized to remove their mean or transformed according to K^{-1} (known). W consists of observations from four views of a moving stereo rig. Let also consider without lack of generality that the planes of the left and right cameras are aligned and that the separation B between the cameras forming the stereo rig defined in one direction (x-axis). In these conditions, the 3D motion $M_i \in \mathbb{SE}(3)$ of the stereo rig will lead to rank$(W) \leq 6$ as expressed in (9).*

$$W = \underbrace{\begin{bmatrix} x_i^{l,1} & \cdots & x_i^{l,N_c} \\ x_i^{r,1} & \cdots & x_i^{r,N_c} \\ x_{i+1}^{l,1} & \cdots & x_{i+1}^{l,N_c} \\ x_{i+1}^{r,1} & \cdots & x_{i+1}^{r,N_c} \end{bmatrix}}_{\substack{8\times N_c \\ \textit{rank}\leq 6}} = h^{-1}\left(\underbrace{\begin{bmatrix} C_i^l \\ C_i^r \\ C_{i+1}^l \\ C_{i+1}^r \end{bmatrix}}_{12\times4} \underbrace{\begin{bmatrix} X_i^{l,1} & \cdots & X_i^{l,N_c} \end{bmatrix}}_{4\times N_c} \right). \tag{9}$$

Here, $h^{-1}(\cdot)$ is a non-linear function that converts each point from homogeneous to non-homogeneous coordinates. This operation rises the maximum rank from 4, the case explained before for 3D data, to 6. In the following section we explain how to perform outlier detection by exploiting Lemma 1.

Proof. A direct geometrical proof can be derived by considering the degrees of freedom of points in a sequence of stereo frames (see Fig. 2). Given an arbitrary point x_i^l seen by the left camera C_i^l, the location of its counterpart x_i^r in

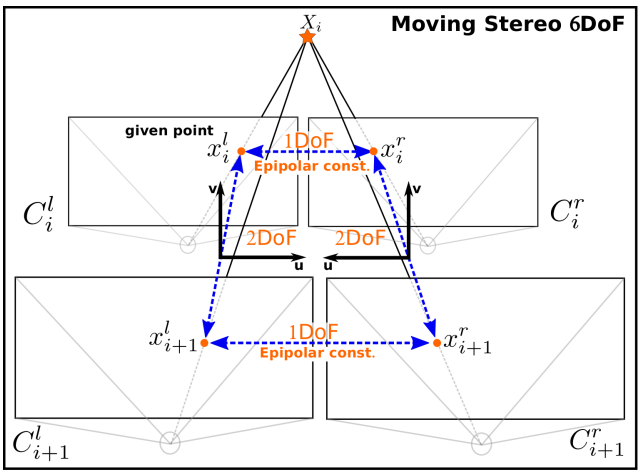

Fig. 2. Diagram representing the geometrical configuration of a moving stereo camera. We show a 3D point X_i projected onto the four views and the arising six degrees of freedom.

C_i^r is determined by an epipolar constraint with one DoF. The same is true for the match between x_{i+1}^l and x_{i+1}^r. Finally the associations along time, $x_i^l \leftrightarrow x_{i+1}^l$ and $x_i^r \leftrightarrow x_{i+1}^r$, given by the motion transformation M^i are determined by two more DoF, summing up to six DoF. □

A. Robust Decomposition with Constrained Rank

We use the rank constraint in (9) to detect potential outliers present in the data matrix W. The corrupted matrix W can me decomposed as a sum of two terms $W = L + S$. In this decomposition we assume that L is the uncorrupted version of W, which satisfies the correct rank constraint. On the other hand, S is a sparse matrix accounting for the outliers, i.e., "infrequent events" of unbounded magnitude. This decomposition can be achieved by imposing sparsity constraints on the error matrix S as

$$\begin{aligned} &\min_{L,S} \ \text{rank}(L) + \lambda \left\|S\right\|_{\ell_0} \\ &\textbf{s.t.} \ \text{rank}(L) \leq r, \left\|W - L - S\right\|_F \leq \epsilon. \end{aligned} \tag{10}$$

However, this formulation becomes intractable due to the presence of the ℓ_0-norm, whose minimization is of combinatorial nature and leads to a NP-hard optimization problem [23] [1] [2]. Recent discoveries have led to the proposal of convex counterparts of (10), referred in the literature to as Robust PCA or Robust Principal Component Pursuit [24]. These techniques are based on a convex relaxation of (10) as

$$\min_{L,S} \left\|L\right\|_* + \lambda \left\|S\right\|_{\ell_1} \text{ s.t. } W = L + S, \tag{11}$$

where $\left\|\cdot\right\|_*$ is the nuclear norm that acts as the convex envelope of the rank function, and the ℓ_1-norm acts as the convex envelope of the ℓ_0 norm.

The convex problem (11), although very convenient, suffers from two conceptual limitations when applied to low-dimensional ambient spaces. First, it does not impose the required rank constraint, i.e., $\text{rank}(L) \leq r$. This could produce low-rank solutions with different ranks, generally smaller, than the one arising in our problem, and thus leading to wrong solutions. Second, and more important, (10) and (11) are not equivalent in those cases where the rank is close to the dimensions of the matrix W [1] [2], as in this case. The behaviour of these limits is abrupt, which is why this phenomenon is known as phase transition, having an universal character, as stated in [4]. For most methods, the region of applicability is approximately of the form

$$\begin{aligned} \rho_r &= \frac{rank(L)}{min(m,n)} \\ \rho_s &= \frac{\left\|S\right\|_0}{mn} \\ \textbf{s.t.} \ \rho_r + \rho_s &\leq k, \end{aligned} \tag{12}$$

where, ρ_r is the rank fraction and ρ_s is the proportion of sparsity (i.e., outliers). The constant k depends on the specific formulation but it is usually $k \approx 0.4$ (see [1] [2] [5]). Also, RPCA convergence theorems assume that the support of S is uniformly distributed, meaning that cases where the sparse components are localized in blocks, rows or columns are out of the applicability of RPCA. Nonetheless, some experiments have been reported in [3] [17], where S has support on columns or rows and is still applicable, as far as the other requirements are met. To this end, [3] proposes the $\ell_{2,1}$-norm as a natural way to deal with outliers as columns. However, in our tests, when the ambient dimension is low, the behaviour of $\ell_{2,1}$ methods is equal or worse than the ℓ_1 counterpart, and computationally more costly.

Our proposal to deal with these problems consists of extending the region of applicability of k by modifying the program (11) to explicitly incorporate the rank constraints. To this end, we reformulate (11) as

$$\begin{aligned} &\min_{L,S} \frac{1}{2} \left\|W - L - S\right\|_F^2 + \lambda \left\|S\right\|_{\ell_1} \\ &\textbf{s.t.} \ \text{rank}(L) \leq r \end{aligned} \quad . \tag{13}$$

We refer to (13) as Robust Decomposition with Constrained Rank (RD-CR). This formulation enables solving problems in harder conditions, i.e., higher ranks and greater proportions of outliers. However, in motion estimation problems, the rank is still too high to achieve an exact estimation of L and S. For this reason, we tackle the problem by using the residual matrix S to infer which columns (point matches) are outliers. The steps required to perform this process are explained in detail in section IV.

IV. Motion Estimation via RD-CR

This section explains our proposal of two stages framework for robust estimation of motion models (see Fig 1). The first stage uses the RC-RD method as a preprocessing step and the second stage applies a compressed regression procedure that exploits data reduction to achieve high performance.

A. RD-CR as a preprocessing stage

The process of outlier detection starts by solving the program (13) with the proposed RD-CR approach, which is based on proximal gradients, following the philosophy of APG [5]. However, in our case the use of the nuclear norm is avoided and the SVD is changed for a more convenient skinny SVD

that serves to generate feasible candidates of rank $r = 6$ with a computational complexity of $O(rn^2)$ per iteration [25]. Our technique also serves to deal with matrices presenting a rank below 6, as the skinny SVD will produce singular values very close to zero in those cases. This method, summarized in Algorithm 1, consisting of three conceptual steps; namely, initialization, projected gradient descent and continuation.

1) ***Initialization*** *(lines 1–3):* The imposed constrained rank makes (13) non-convex, for which a proper initialization of L_0 and S_0 is required. A good initialization is achieved by solving the convex surrogate (11) through the use of the APG method [5]. APG uses a standard proximal algorithm improved by Nesterov acceleration [26] and a continuation technique to speed up the convergence of the terms involving proximity operators [27], i.e. both the nuclear and ℓ_1 norms. Despite its convexity, it should be noticed that in the extreme cases we are addressing (high ρ_r) APG might not converge to a suitable solution. However, when APG is run for a fixed amount of iterations ($k = 20$), it produces candidates with enough quality to initialize our method. The remaining parameters are set to $\lambda = 10^{-2}$, $\bar{\mu} = 10^{-9}$ and $k_{\max} = 20$.

2) ***Projected Gradient Descent*** *(lines 6–13):* The main loop of Algorithm 1 performs $k_{\max}$ iterations, each one consisting of the computation of a projected descent step for L_k and S_k. To this end, descent steps G_k^L and G_k^S are computed from the Euclidean gradient $D_k = L_k + S_k - W$. G_k^L is projected to the rank-6 manifold through the skinny SVD, while G_k^S is projected to promote sparsity through the soft-thresholding operator

$$\mathcal{S}_\mu[M] = \max(0, M - \mu) + \min(0, M + \mu). \tag{14}$$

We use fixed-length steps α_L, α_S for the descent process with values in the ranges $\alpha_L \in [1/10, 1]$, $\alpha_S \in [1/10, 1/2]$, which are close to the Lipschitz constant of (13).

3) ***Continuation*** *(lines 4 and 14–16):* Continuation is defined as a technique to improve the convergence of proximity operators like the soft-thresholding by dynamically adapting the parameter μ_k following a geometrical series [27] [5]. However, in our approach we adapt the continuation scheme to be more coherent with the current estimation. This is done by taking μ_k as a fraction of the mean size of the singular values of the Euclidean gradient D_k, as,

$$\mu_D = \delta \left\| D_k \right\|_F / \sqrt{mn}, \quad \delta = 10^{-3}. \tag{15}$$

Then, μ_D is used to generate a new $\mu_{k+1} = \max(\frac{1}{\lambda}\mu_D, \bar{\mu})$, which empirically led to a higher performance.

A proof of convergence for the Algorithm 1 is very hard due to the large value of ρ_r for our low ambient dimension scenario. However, this does not limit its applicability to real problems since in our experiments running the algorithm during $k_{\max} = 20$ iterations has produced good results. Moreover, the proposed method is very efficient due to the moderate computational complexity of Algorithm 1, i.e. $O(rn^2)$, and given that $r = 6$ and $n = 8$.

After applying Algorithm 1, the resultant S contains outliers distributed in columns. Then, in order to decide if the j-column (point match) is an outlier or not, we apply the following decision criterion,

Algorithm 1 RD-CR with Proximal Gradients

Require: Data matrix $W \in \mathbb{R}^{m\times n}$, initialization $(L^{\text{APG}}, S^{\text{APG}})$, λ, $k_{\max}$, δ, $\bar{\mu}$.
1: $k \leftarrow 0$
2: $L_k \leftarrow L^{\text{APG}}$, $S_k \leftarrow S^{\text{APG}}$, $\alpha_L \leftarrow 1$, $\alpha_S \leftarrow 1/5$
3: $(U_0, \Sigma_0, V_0) \leftarrow svds(W, 6)$, $\tilde{W} \leftarrow U_0\Sigma_0 V_0^T$
4: $\mu_k \leftarrow \delta \left\| W - \tilde{W} \right\|_F / \sqrt{mn}$.
5: **while** $k < k_{\max}$ **do**
6: $D_k \leftarrow L_k + S_k - W$
7: *//Update of L as a descent gradient step + projection*
8: $G_k^L \leftarrow L_k - \alpha_L D_k$.
9: $(U_{k+1}, \Sigma_{k+1}, V_{k+1}) \leftarrow svds(G_k^L, 6)$
10: $L_{k+1} = U_{k+1}\Sigma_{k+1}V_{k+1}^T$.
11: *//Update of S as a descent gradient step + projection*
12: $G_k^S \leftarrow S_k - \alpha_S D_k$.
13: $S_{k+1} \leftarrow \mathcal{S}_{\mu_k}[G_k^S]$.
14: *//Update of μ_k with continuation*
15: $\mu_D \leftarrow \left\| D_k \right\|_F / \sqrt{mn}$
16: $\mu_{k+1} \leftarrow \max(\frac{1}{\lambda}\mu_D, \bar{\mu})$.
17: $k \leftarrow k + 1$.
18: **end while**
19: **return** $L \leftarrow L_k$, $S \leftarrow S_k$.

$$\psi_S(x_j) : \mathbb{R}^8 \rightarrow \{0, 1\} \tag{16}$$

$$\psi_S(x) = \begin{cases} 1, & \text{if } \|x\|_{\ell_1} > \min(\tau_0, \frac{1}{N_c}\|S\|_{\ell_1}) \\ 0, & \text{otherwise.} \end{cases} \tag{17}$$

The decision made by $\psi_S(\cdot)$ is based on a threshold directly proportional to the average ℓ_1-norm of the matrix S up to a saturation point $\tau_0 = 0.5$. This expression serves to adapt the threshold to the mean residual of the data, what in our experiments has reported the best performance through cross-validation. In the cases where $\psi_S(x_j) = 1$, the $j - th$ match point is removed from the correspondence set.

B. Fast Compressed Least-Squares

The second stage of our framework consists in the actual estimation of the transformation M_i. In order to make this stage as fast as possible we consider those cost functions that are easy to linearise with respect to the unknowns. We have selected the algebraical cost function proposed in [7], defined as;

$$E(M) = \sum_{j=1}^{N} \left\| \left(\Pi\text{K } M\hat{X}_i^{l,(j)} \right) \times \hat{x}_{i+1}^{l,(j)} \right\|^2 + \left\| \left(\Pi\text{K } M\hat{X}_i^{l,(j)} - \vec{B} \right) \times \hat{x}_{i+1}^{r,(j)} \right\|^2. \tag{18}$$

This function measures the reprojection error of points in the two stereo views as the misalignment of homogeneous vectors. Since this cost is algebraical instead of a geometric one, it requires a proper normalization of the data, i.e., removing its mean and scaling it up by its standard deviation as described in [6] [28]. In this context, $\hat{X}$ stands for homogeneous 3D points and $\hat{x}$ for homogeneous 2D points. $\times$ is the standard

cross product of vectors. The rest of the notation remains as explained before.

The choice of (18) is motivated by its tolerance to be efficiently compressed. Compression consists in turning (18), which sums over all the observations, into a compact equivalent Γ, known as Reduced Measurement Matrix (RMM). In this case, the compression technique used is fully equivalent to the original data, i.e., is a lossless compression. The compression step is performed by following the process,

$$\begin{aligned} E(M) = \sum_{j=1}^{N} \left\| \left(\Pi \mathrm{K}\ M \hat{X}_i^{l,(j)} \right) \times \hat{x}_{i+1}^{l,(j)} \right\|_{\ell_2}^2 + \\ \left\| \left(\Pi \mathrm{K}\ M \hat{X}_i^{l,(j)} - \vec{B} \right) \times \hat{x}_{i+1}^{r,(j)} \right\|_{\ell_2}^2 = \\ \sum_{j=1}^{N} \left\| A^{l,(j)} \breve{M} \right\|_{\ell_2}^2 + \left\| A^{r,(j)} \breve{M} \right\|_{\ell_2}^2 = \\ \left\| \boldsymbol{A^l} \breve{M} \right\|_{\ell_2}^2 + \left\| \boldsymbol{A^r} \breve{M} \right\|_{\ell_2}^2 = \\ \breve{M}^T \underbrace{\boldsymbol{A^l}^T \boldsymbol{A^l}}_{\boldsymbol{\Gamma^l}} \breve{M} + \breve{M}^T \underbrace{\boldsymbol{A^r}^T \boldsymbol{A^r}}_{\boldsymbol{\Gamma^r}} \breve{M} = \breve{M}^T \boldsymbol{\Gamma} \breve{M}. \end{aligned} \quad (19)$$

For our particular problem Γ is a 13×13 matrix, which is determined from the number of coefficients required to linearise the original cost function (18). Here, $A^{m,(j)}$, for $m = l, r$ is the matrix of linearised coefficients which results from combining the calibration K, Π and the observations, as shown in (24). Then $\boldsymbol{A^m}$ is the stack of the N matrices $A^{m,(j)}$ and $\breve{M} = [\mathbf{stack}(R), T, 1]^T$, i.e., a vector resulting from stacking the components of M with an homogeneous component.

The reduction of (19) is very efficient due to its algebraical nature and just has to be performed once, before the optimization takes place. Then, the optimization is posed as a compressed least-squares problem on the $\mathbb{SE}(3)$ manifold in order to fulfil the orthonormal constraint of 3D rigid motions. This can be formulated as,

$$\begin{aligned} \omega^* &= \underset{\omega}{\operatorname{argmin}}\, E(M) = \underset{\omega}{\operatorname{argmin}}\, \breve{M}^T \boldsymbol{\Gamma} \breve{M} \\ &= \underset{\omega}{\operatorname{argmin}}\, \mathbf{stack}[\exp_{\mathbb{SE}(3)}(\omega)]^T\ \boldsymbol{\Gamma}\ \mathbf{stack}[\exp_{\mathbb{SE}(3)}(\omega)], \end{aligned} \quad (20)$$

where ω^* is the estimated motion transformation given as an element of the Lie algebra, $\omega \in \mathfrak{se}(3)$, which is the tangent space of $\mathbb{SE}(3)$ at the identity. The use of the exponential map serves to move from the tangent space back to the manifold. This map is given by the Rodrigues formula [29] and is a requirement to perform the evaluation of the cost function. Despite the special properties of (20), it can be easily optimized by using Levenberg-Marquardt. The only requirement for this optimization is the computation of the Jacobian of (20), something that is done by computing its partial derivatives as,

$$\begin{aligned} \partial_{\omega_i} E(M) \partial_{\omega_i} \xi^T (\boldsymbol{\Gamma} + \boldsymbol{\Gamma}^T) \xi, \\ \text{where } \xi = \mathbf{stack}[\exp_{\mathbb{SE}(3)}(\omega)] = \breve{M}. \end{aligned} \quad (21)$$

The partial derivatives of $\exp_{\mathbb{SE}(3)}$ are calculated by considering a k-order approximation (with $k \approx 10$ in our experiments) of the expression from the Taylor's series defined as,

$$\begin{aligned} \partial_{\omega_i} \mathbf{exp}(Z) \approx \partial_{\omega_i} \left(\sum_{n \geq 0}^{k} \frac{Z^n}{n!} \right) = \\ \sum_{n \geq 1}^{k} \frac{1}{n!} \left(\partial_{\omega_i} Z\ Z^{n-1} + Z\ \partial_{\omega_i} Z^{n-1} \right) \end{aligned} \quad (22)$$

$$\text{with } Z = \begin{bmatrix} 0 & -\omega_3 & -\omega_2 & \omega_4 \\ \omega_3 & 0 & -\omega_1 & \omega_5 \\ -\omega_2 & \omega_1 & 0 & \omega_6 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (23)$$

The optimization of this function is very fast and converges in a couple of iterations. The use of Γ reduces each iteration to just various matrix by vector products and vector by vector products, what makes the compressed least-squares procedure very efficient. The total time for the compression and the optimization is around one millisecond in a standard desktop computer.

V. Experimental Evaluation

This section shows the experiments performed on synthetic and real data to validate our approach. On one hand, synthetic data is used as a control environment to analyse the behaviour of the RD-CR algorithm in different situations. On the other hand, experiments with real data serve to compare the performance of the whole framework against state-of-the-art techniques for stereo Visual Odometry. From the results we conclude that our approach is competitive against state-of-the-art methods in terms of accuracy and is more efficient in terms of computation.

A. Synthetic data tests

We have created a set of tests that simulate the conditions of a stereo rig performing a rigid motion in 3D, which is the problem of application addressed in this paper. For each case, we evaluate the method for a specific matrix size N_c (i.e., number of matches) and proportion of outliers $P_o = \rho_s$. We define the stereo rig with the reference camera matrices (C_i^l, C_i^r). This means that C_i^l is the reference frame, while C_i^r is set at baseline distance B. Then, $N_m > 50$ motion transformations

$$A^{m,(j)} = \begin{bmatrix} [0]_{1\times 3} & f X_i^{l,(j)^T} & (c_v - v_{i+1}^{m,(j)}) X_i^{l,(j)^T} & 0 & f & (c_v - v_{i+1}^{m,(j)}) & 0 \\ -f X_i^{l,(j)^T} & [0]_{1\times 3} & (u_{i+1}^{m,(j)} - c_u) X_i^{l,(j)^T} & -f & 0 & (u_{i+1}^{m,(j)} - c_u) & \alpha \\ f v_{i+1}^{m,(j)} X_i^{l,(j)^T} & -f u_{i+1}^{m,(j)} X_i^{l,(j)^T} & (c_u v_{i+1}^{m,(j)} - c_v u_{i+1}^{m,(j)}) X_i^{l,(j)^T} & f v_{i+1}^{m,(j)} & -f u_{i+1}^{m,(j)} & (c_u v^{(j)} - c_v u_{i+1}^{m,(j)}) & \beta \end{bmatrix} \quad (24)$$

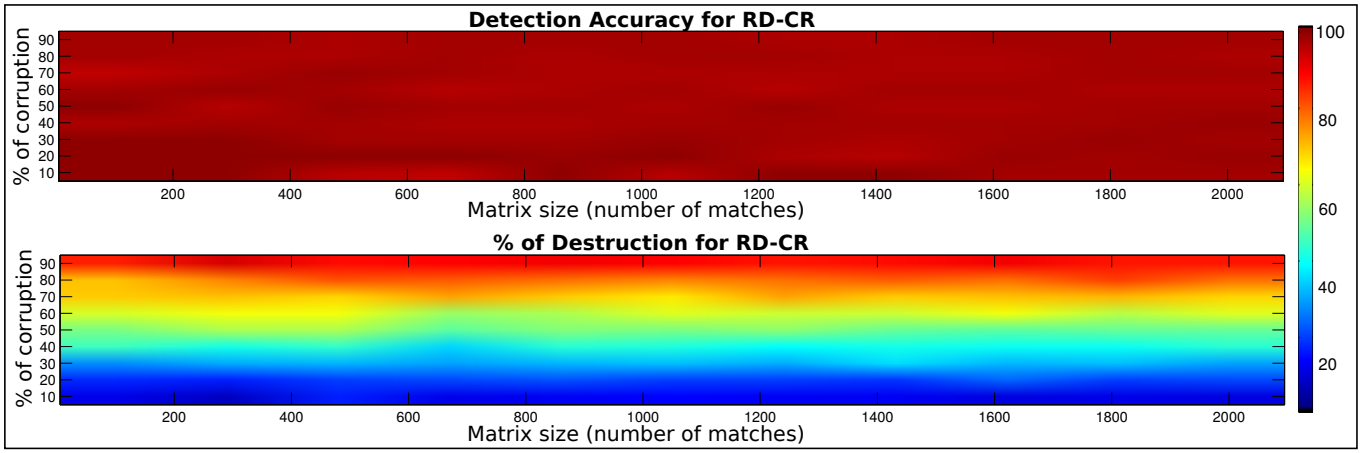
Detection Accuracy for RD-CR
% of corruption
Matrix size (number of matches)
% of Destruction for RD-CR
% of corruption
Matrix size (number of matches)

Fig. 3. **Top**, mean accuracy evaluation of RD-CR w.r.t. the number of matches and the percentage of outliers. **Bottom**, mean percentage of elimination made by RD-CR for each case.

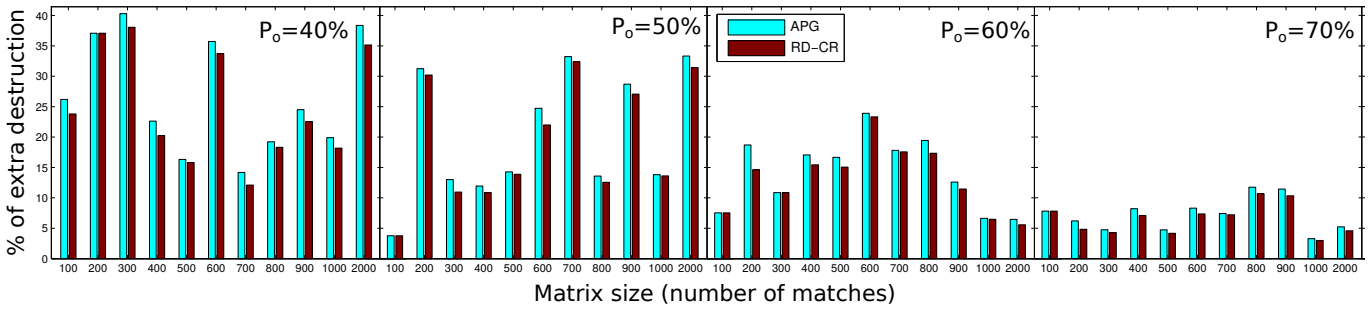

Fig. 4. Comparison of APG and RD-CR regarding the percentage of excessive elimination of outliers.

M_i are created such that each $M_i \in \mathbb{SE}(3)$. For each model M_i we create N_c matches $\{x_i^l, x_i^r, x_{i+1}^l, x_{i+1}^r\}$, according to the four views. Point matches are then contaminated with random noise of magnitude $\sigma_n = 1.5$ and impulsive noise with magnitude $\sigma_j \in [2, 100]$, to represent outliers. In each case the $P_o\%$ of elements $x_i^l, x_i^r, x_{i+1}^l, x_{i+1}^r$ are corrupted by σ_j in all their components. This data is reorganized to form a matrix $W_i \in \mathbb{R}^{8\times N_c}$ as explained in Sec. III.

We have run this test repeatedly for values of N_c from 100 to 2000 matches and ratios of outliers P_o from 10% to 90%. In each combination of N_c and P_o, 50 repetitions are performed and from them the average of outlier detection is calculated, i.e., the accuracy of the method. The results of this experiment are shown in Fig. 3**(Top)**. Our approach RD-CR shows a high success ratio, identified by warm colours, for the different sizes of matches considered and percentages of outliers. Additionally, as we can see in Fig. 3**(Bottom)**, the percentage of outlier removal of RD-CR closely matches the expected percentage of corruption, meaning that the method recall is high and it does not produce many false positives.

By imposing the correct rank, RD-CR reduces the percentage of excessive elimination of outliers, a problem that is present in APG, as shown in Fig. 4. Given several levels of corruption we compare the excess of outlier removal (false positives) caused by RD-CR and APG. From these results we can conclude that, in general, RD-CR overcomes the efficacy of APG to detect outliers and causes less false positives for the different configurations of noise. This is mainly due to the correct rank estimation done by RD-CR, as APG tends to underestimate the rank, leading to more false positives.

This synthetic experiment aims to show that it is possible to exploit RPCA-like methods to detect outliers even when working beyond the recovery boundaries of these techniques, which is the initial assumption of this work. In these cases, the high accuracy of our approach at outlier detection is closely linked to a precise selection of the threshold τ_0 within the function ψ_S. However, when working in real scenarios this information cannot be determined in a precise fashion. In addition, real scenarios bring extra challenges, such as multiple corruption per correspondence, i.e., more than one point is severely affected within a match, and σ_j can adopt moderated values, making the detection process more challenging. Nevertheless, our next experiments prove that RD-CR still work in real scenarios under these adverse conditions.

B. Real data tests

We now show the good performance of our framework applied to stereo Visual Odometry in real sequences. Data sets are taken from the challenging KITTI benchmark [8] in its Visual Odometry modality. KITTI sequences correspond to large urban environments, providing stereo images taken from a car. We use sequences 00–10 for which ground truth is available.

The comparison is carried out with 5 state-of-the-art methods for motion estimation: RANSAC [11] and L1-averaging [13] representing classical consensus and modern

consensus-by-averaging methods. IRLS-L1 [16] and Huber M-estimator [14] as representatives of fast robust estimators, and Compressed Least-squares (C-LS) —the compressed flavour of least-squares presented in section IV-B— stands for non-robust methods. In addition, APG is also tested in order to highlight the difference in accuracy when the rank constraint is not imposed. Other methods, like Bundle Adjustment, are excluded since our goal is the estimation of individual models, without considering the temporal relationships among the different motion transformations.

In our experiments RANSAC and L1-averaging are set to produce a maximum of 250 and 1000 models respectively during the consensus, which is a realistic number for real-time applications. The constant for the Huber M-estimator is set to $\delta = 2.5$, after a process of validation. The number of point matches per frame is around 2000. For all methods the computation of the models is done by an optimization process on $\mathbb{SE}(3)$, via a couple of iterations of Levenberg–Marquardt. The cost function used in this optimization follows a standard projection of 3D points into the left and right cameras [19], over a fixed number of matches $N_k = 3$, i.e., the minimum amount required to generate a valid model. In our experiments this cost function is defined as,

$$\underset{\omega}{\operatorname{argmin}} \sum_{i=1}^{N_k} \left\| \Pi \mathrm{K} \exp(\omega) X_i^{l,(j)} - x_{i+1}^{l,(j)} \right\|_{\ell_2}^2 + \left\| \Pi \left(\mathrm{K} \exp(\omega) X_i^{l,(j)} - \vec{B} \right) - x_{i+1}^{r,(j)} \right\|_{\ell_2}^2 . \tag{25}$$

As an error metric we use the mean of the individual relative errors along a whole sequence. To measure this error, we have considered the most natural metric, i.e., the group structure of rigid transformations, which is defined as;

$$\mathcal{E}(M_i, M^*) = \left\| log_{\mathbb{SE}(3)}(M_i {M^*}^{-1}) \right\|_{\ell_2} . \tag{26}$$

In this expression, $log_{\mathbb{SE}(3)}(\cdot)$ is the logarithm of the $\mathbb{SE}(3)$ group, mapping from the group to its tangent space. This tool is used to measure the difference between the estimated motion M_i and the ground truth M^*. From that, a relative error is computed to facilitate an intuitive interpretation,

$$\mathcal{E}_{\text{rel}}(M_i, M^*) = \frac{\left\| log_{\mathbb{SE}(3)}(M_i {M^*}^{-1}) \right\|_{\ell_2}}{\left\| log_{\mathbb{SE}(3)}(M^*) \right\|_{\ell_2} + \epsilon}, \text{ for } \epsilon = 10^{-5}. \tag{27}$$

It should be stressed that this metric provides a measure of the local error, and does not take into account the accumulative effect of the error in the global trajectories. This point is critical in order to understand the differences between the results provided by Table I, which is focused on local errors; and Fig. 5 and Fig. 6, which show the accumulative effects of the errors.

Table I summarizes the average results of the selected methods when applied to the KITTI sequences for the task of Visual Odometry. The analysis of this table reveals that RD-CR leads to results at the level of the consensus methods. Moreover, our proposal is more efficient since it does not require to generate multiple hypotheses, being able to detect outliers from the structure of the data and leading to a promising trade-off (see Table II).

As shown, RD-CR produces better results than the L1-averaging for all the sequences but seq. 01. RD-CR also performs similar to RANSAC, achieving better results for sequences 02, 05 and 07. In general, our approach outperforms robust regression methods like IRLS-L1 and the Huber M-estimator. In challenging scenes, such as sequence 01, 07 or 10, RD-CR produces significantly better results, while the IRLS-L1, Huber method and C-LS are seriously affected by outliers.

Our results also show that skipping the rank constraint has a dramatic negative influence in the outlier detection process. For this reason, APG performs dramatically worse than the proposed RD-CR technique. There are two phenomena involved in this problem. When the rank is overestimated, APG will miss outliers (false negatives), since they will be absorbed into the low-rank matrix L. On the contrary, if APG underestimates the rank, the number of linear dependencies in L will increase and some elements will migrate to S (false positives), losing some inliers needed for the good estimation of the solution.

Figure 5 shows results from the point of view of the estimated trajectories. The analysis of the first two plots, Fig. 5**(a)**–**(b)** shows that the accuracy of our proposal is at the level of RANSAC or even higher. This is depicted in the zoomed region of Fig. 5**(b)**, where the trajectory estimated by RD-CR is closer to the ground truth. This behaviour is repeated for the trajectories shown in Fig. 5**(c)**,**(d)**,**(e)** and **(f)**, where the difference of RD-CR with respect to the others is more pronounced. In the last sequences **(c)**–**(f)**, the trajectories estimated by RD-CR are several meters closer to the ground truth than those from the state-of-the-art methods.

Finally, we would like to consider the failure cases. In the case of seq. 01, depicted in Fig. 6**(a)**, we attribute the bad results to the severe visual aliasing found in this highway scene, a phenomenon leading to a very high ratio of matches that are outliers but behave similarly to inliers, i.e., they are mistaken for a valid $\mathbb{SE}(3)$ motion. It is also fair to highlight that robust estimator like IRLS-L1 perform very well in some sequences. For instance, in seq. 04 (see Fig. 6**(b)**), IRLS-L1 produces excellent results. The main problem RD-CR finds in this sequence is the matches corresponding to a car moving at the same speed that our vehicle. This subtle motion is wrongly combined with the predominant motion transformation leading to a small but continuous drift through the sequence. We are currently working on a way to solve this specific issue for future versions of the approach.

1) Complexity evaluation: We present the excellent computational performance of the proposed approach by comparing it against selected state-of-the-art techniques. A special effort has been made in order to produce a fair comparison of the methods. Since all the algorithms are coded in Matlab one needs to account for the bad scalability of loops, which usually lead to a super-linear cost with respect to the loop length. To avoid this problems we analyse iterative blocks in terms of their average number of iterations and the average cost of each iteration. The average cost per iteration is calculated by

TABLE I
EVALUATION OF METHODS MEASURING THE MEAN RELATIVE ERROR FOR KITTI SEQUENCES (00-10).

	00	01	02	03	04	05	06	07	08	09	10	**Global**
Consensus based Methods												
RANSAC	5.27%	10.05%	4.83%	5.97%	5.30%	8.47%	5.52%	13.89%	7.77%	5.45%	5.80%	6.71%
L1-Avg.	5.44%	19.55%	4.95%	6.34%	5.72%	8.23%	5.92%	13.08%	8.26%	5.84%	6.79%	7.35%
Non-Consensus based Methods												
IRLS-L1	5.35%	30.68%	4.86%	6.08%	5.44%	7.66%	5.63%	13.22%	7.94%	5.57%	7.35%	7.71%
Huber-M.	5.30%	34.40%	4.80%	6.10%	9.20%	11.89%	5.59%	25.61%	12.95%	6.64%	7.97%	9.88%
C-LS	5.77%	51.52%	5.05%	7.23%	6.45%	40.99%	6.32%	135.32%	10.70%	6.07%	8.10%	19.22%
Robust Decomposition based Methods												
APG	10.19%	22.78%	7.59%	9.11%	7.27%	16.74%	6.67%	24.14%	13.67%	10.49%	10.64%	12.12%
RD-CR (Ours)	5.42%	27.23%	4.76%	6.02%	5.69%	7.85%	5.36%	12.15%	8.00%	5.46%	6.32%	7.49%

TABLE II
COMPUTATIONAL TIME ANALYSIS FOR THE METHODS UNDER COMPARISON. THE EVALUATION IS CARRIED OUT IN MATLAB BY MEASURING THE DIFFERENT PARTS OF EACH METHOD INDEPENDENTLY AND KEEPING THE SIZE OF THE LOOPS SMALL TO PROVIDE FAITHFUL COMPARISONS.

Methods	**Time analysis**		**Time (s)**
RANSAC	$N_{\text{model}} \times (T_{\text{ModelGeneration}} + N_{\text{points}} \times T_{\text{CheckInliers}})$	$250 \times (1.3e{-}3 + 2000 \times 5e{-}5)$	25.32
L1-Avg	$N_{\text{model}} \times T_{\text{ModelGeneration}} + T_{\text{InitL1}} +$ Iters$\times (N_{\text{model}} \times T_{\text{ProcessModel}} + T_{\text{RetractionMap}})$	$1000 \times 1.3e{-}3 + 1.7e{-}2 +$ $200 \times (1000 \times 7.26e{-}5 + 4.7e{-}5)$	15.85
IRLS-L1	Iters$\times N_{\text{points}} \times (T_{\text{CostJacobian}} + T_{\text{Reweighting}})$	$15 \times 2000 \times (1.2e{-}4 + 2.8e{-}5)$	4.44
Huber-M	Iters$\times N_{\text{points}} \times T_{\text{CostJacobian}}$	$15 \times 2000 \times 1.2{-}4$	3.60
C-LS	$T_{\text{CreationRMM}} +$ Iters$\times T_{\text{CostJacobian}} + T_{\text{Refinement}}$	$1.6e{-}2 + 5 \times 6.2e{-}5 + 2.6e{-}1$	0.28
APG	Iters$\times (T_{\text{Gradients}} + T_{\text{SVD-ECON}} + T_{\text{Nuclear}} + T_{\text{Proximity}} + T_{\text{Nesterov}}) + T_{\text{C-LS}}$	$100 \times (1.2e{-}3 + 1.37e{-}3 + 3e{-}3 + 7.13e{-}4 + 8.75e{-}4) + 0.28$	0.99
RD-CR (Ours)	ItersAPG$\times (T_{\text{Gradients}} + T_{\text{SVD-ECON}} + T_{\text{Nuclear}} + T_{\text{Proximity}} + T_{\text{Nesterov}}) +$ ItersRDCR$\times (T_{\text{Gradients}} + T_{\text{SSVD}} + T_{\text{Proximity}}) + T_{\text{C-LS}}$	$20 \times (1.2e{-}3 + 1.37e{-}3 + 3e{-}3 + 7.13e{-}4 + 8.75e{-}4) +$ $20 \times (2.64e{-}4 + 1.4e{-}3 + 7.13e{-}4) + 0.28$	0.47

TABLE III
TIME LEGEND, LISTING THE DIFFERENT COMPONENTS OF EACH ALGORITHM.

N_{model}	Number of model generated from random matches.	$T_{\text{ModelGeneration}}$	Time used to generate each model by means of an iterative process on $\mathbb{SE}(3)$.
N_{points}	The amount of point matches considered.	$T_{\text{CheckInliers}}$	Time employed to check the membership of a specific match to the inlier set.
T_{InitL1}	Time taken by the initialization of the ℓ_1-averaging.	*Iters*	Average number of iterations of a given method
$T_{\text{ProcessModel}}$	Time for the projection of a model to the tangent space $\mathfrak{se}(3)$.	$T_{\text{RetractionMap}}$	Time needed to map a model from its tangent space to the $\mathbb{SE}(3)$ manifold.
$T_{\text{CostJacobian}}$	Time used to compute a cost function and its partial derivatives.	$T_{\text{Reweighting}}$	Time used for updating the weights of $IRLS-L1$ after each iteration.
$T_{\text{CreationRMM}}$	Time for compressing the initial matches into the Reduced Measurement Matrix.	$T_{\text{Refinement}}$	Time used to perform an optional refinement based on the re-estimation of the inlier set.
ItersAPG	Number of iterations needed by APG.	*ItersRDCR*	Number of iterations needed by RD-CR.
$T_{\text{SVD-ECON}}$	Time to compute an economic SVD.	T_{SSVD}	Time to compute a skinny SVD (compact SVD plus thresholding for rank $= 6$).
$T_{\text{Gradients}}$	Time needed to compute the Riemannian gradients.	T_{Nuclear}	Time to compute the nuclear norm.
$T_{\text{Proximity}}$	Time to compute the soft-thresholding operator.	T_{Nesterov}	Time used to perform the computation of approximate points by following Nesterov's rule.
$T_{\text{C-LS}}$	Time required to perform the Compressed Least-squares method.		

using short loops, to avoid the region of bad scalability.

The results of our analysis are summarized in Table II. For each method we highlight the most important blocks that are involved in the computational performance of the approaches.

In this evaluation, RD-CR shows to be one of the most efficient approaches, just behind the non-robust C-LS, which is the second stage of RD-CR. Consensus models spend a large amount of time at generating putative models and checking the pertinence of matches to the inlier set. This is specially noticeable in our experiments since the number of matches is large ($N_{\text{points}} \approx 2000$). Such a choice is justified by works like [7], showing that managing large amounts of matches improves stability and accuracy (up to a limit). Methods based on least-squares, such as the IRLS and the Huber M-estimator suffer from a similar problem due to the computation of large Jacobians.

The behaviour of APG and RD-CR is quite different to the previous methods. These approaches do not deal with individual instances of the data, but with the whole data set as a single entity. This conception allows APG and RD-CR to be very efficient even when having to face a large number of matches. Furthermore, due to the structure of RD-CR, the number of iterations required to its convergence is less than the iterations needed by APG; that is, RD-CR achieves a higher accuracy with less number of iterations, being computationally very appealing.

VI. CONCLUSION AND FUTURE WORK

We have proven that outlier detection in motion estimation problems can benefit from the use of robust decomposition techniques, even for those cases which are beyond the theoretical limits of the robust recovery. We have exploited the information available in deficient recoveries to infer which columns are outliers. To this end, we have proposed a new method called RD-CR, able to impose rank constraints arising in motion estimation problems, in an efficient way. Additionally, we presented a robust regression framework that combines RD-CR as a preprocessing stage, prior to a fast

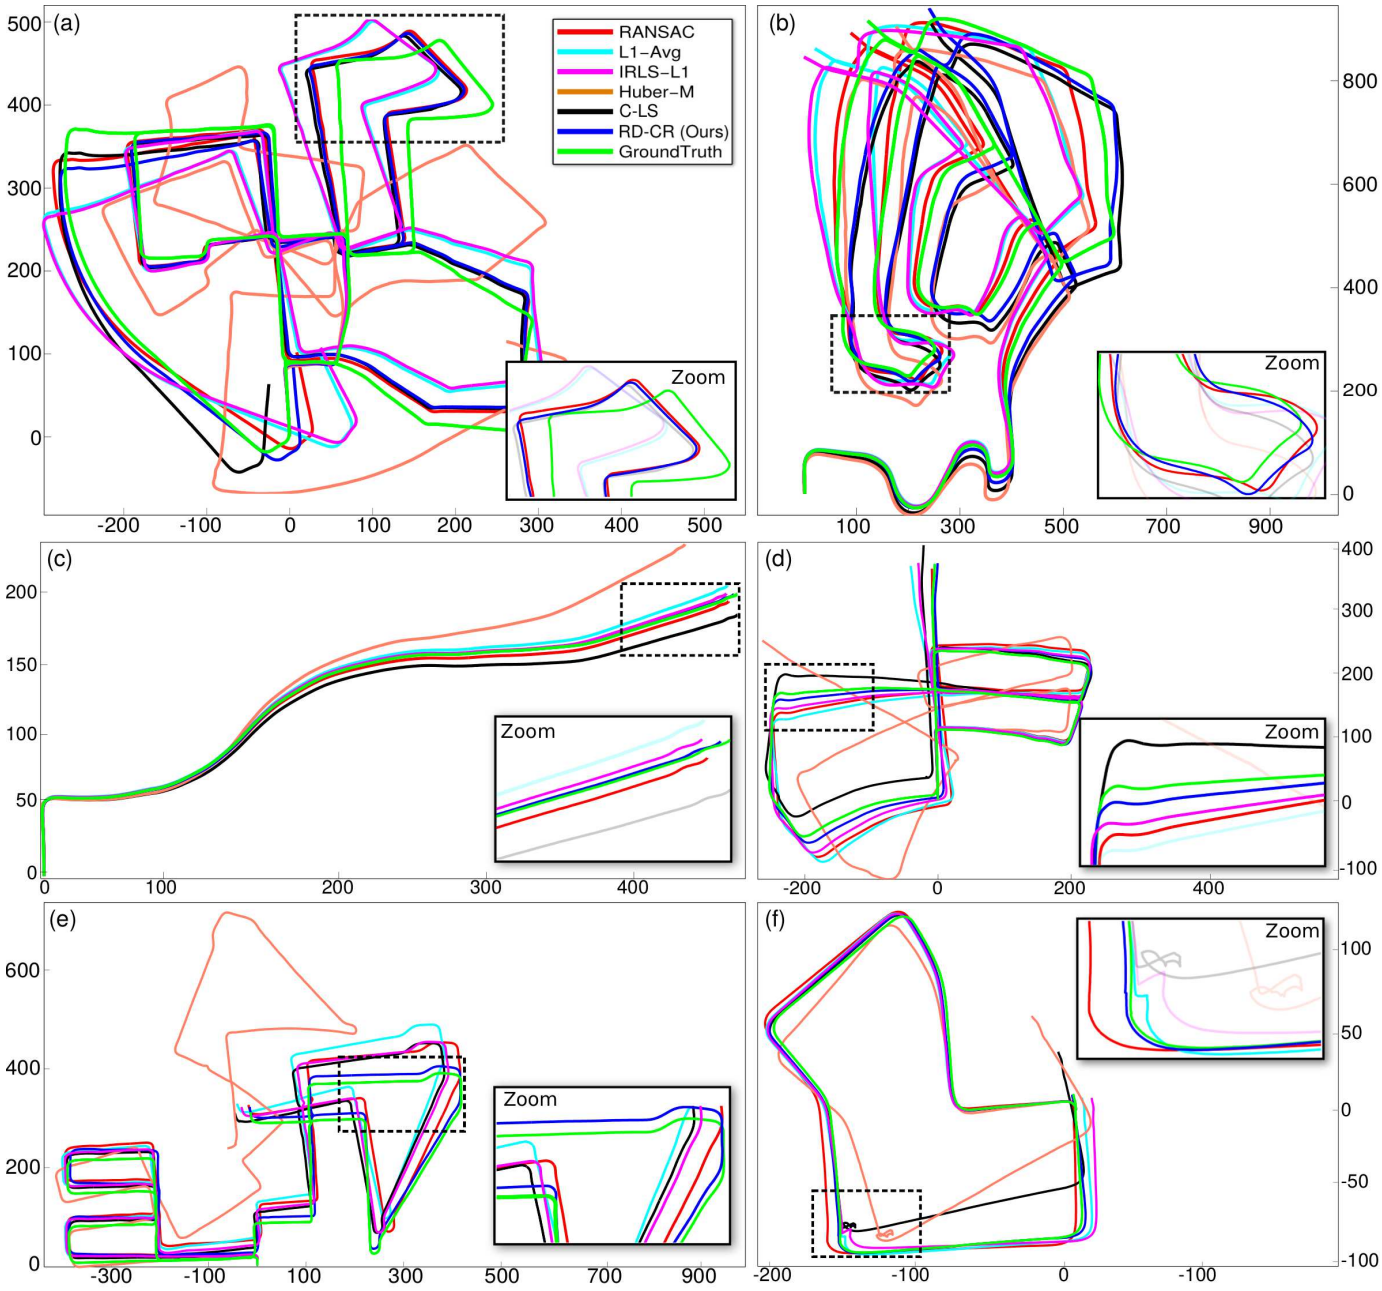

Fig. 5. Visual Odometry results for sequences 00 **(a)**, 02 **(b)**, 03 **(c)**, 05 **(d)**, 08 **(e)** and 07 **(f)**.

compressed least-squares algorithm based on data reduction. We verified that this conception alleviates known problems of compressed least-squares methods and leads to a robust estimator that is both precise and fast. We have validated our approach by performing tests in both synthetic and real data sets, providing a good view of the theoretical behaviour of the method and its good performance in real scenarios. As future improvements, we are experimenting with new ways of addressing rank constraints that have proven to be computationally more efficient and capable of dealing with large-scale problems.

Acknowledgment

This work has been supported by the Universitat Autònoma de Barcelona, the Fundación Séneca 08814PI08, the Spanish government, by the projects FIS201129813C0201; TRA201129454C0301 (eCo-DRIVERS) and NICTA. NICTA is an entity funded by the Australian Government as represented by the Department of Broadband, Communications, and the Digital Economy, and the ARC through the ICT Centre of Excellence Program.

References

[1] J. Wright, "Robust principal component analysis: Exact recovery of corrupted low-rank matrices via convex optimization," in *NIPS*, 2009.

[2] E. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *Journal of the ACM*, vol. 58, no. 3, May 2011.

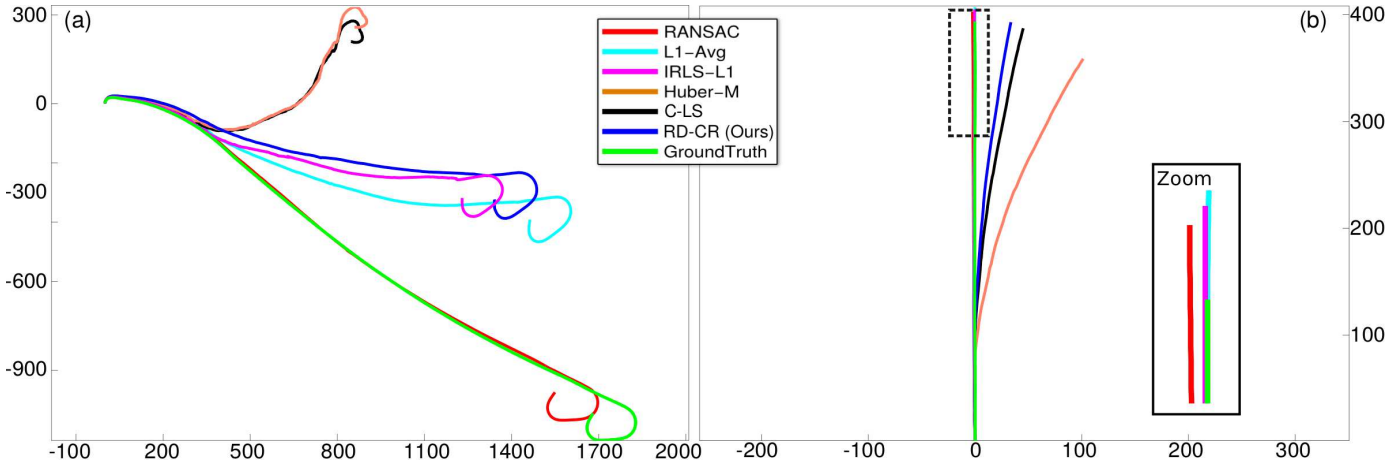

Fig. 6. Visual Odometry results for sequences 01 **(a)**, and 04 **(b)**.

[3] H. Xu, C. Caramanis, and S. Sanghavi, "Robust pca via outlier pursuit," in *NIPS*, 2010.

[4] D. L. Donoho and J. Tanner, "Observed universality of phase transitions in high-dimensional geometry, with implications for modern data analysis and signal processing." *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, vol. 367, Nov. 2009.

[5] Z. Lin, A. Ganesh, J. Wright, L. Wu, M. Chen, and Y. Ma, "Fast convex optimization algorithms for exact recovery of a corrupted low-rank matrix," in *Workshop on Comp. Adv. in Multi-Sensor Adapt. Processing*, 2009.

[6] R. Hartley, "Minimizing algebraic error in geometric estimation problems," in *ICCV*, 1998.

[7] G. Ros, J. Guerrero, A. Sappa, D. Ponsa, and A. López, "Fast and robust l1-averaging-based pose estimation for driving scenarios," in *BMVC*, 2013.

[8] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the KITTI vision benchmark suite," in *CVPR*, 2012.

[9] R. Raguram, J. Frahm, and M. Pollefeys, "A comparative analysis of RANSAC techniques leading to adaptive real-time random sample consensus," in *ECCV*, 2008.

[10] F. R. Hampel, E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel, *Robust Statistics: The Approach Based on Influence Functions (Wiley Series in Probability and Statistics)*. New York: Wiley-Interscience, 2005.

[11] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, 1981.

[12] V. M. Govindu, "Lie-algebraic averaging for globally consistent motion estimation," in *CVPR*, 2004.

[13] R. Hartley, K. Aftab, and J. Trumpf, "L1 rotation averaging using the Weiszfeld algorithm," in *CVPR*, 2011.

[14] P. J. Huber, *Robust Statistics*, ser. Wiley Series in Probability and Statistics - Applied Probability and Statistics Section Series. Wiley, 2004.

[15] P. W. Holland and R. E. Welsch, "Robust regression using iteratively reweighted Least-Squares," *Communications in Statistics: Theory and Methods*, pp. 813–827, 1977.

[16] D. P. O'Leary, "Robust regression computation using iteratively reweighted least squares," *SIAM Journal on Matrix Analysis and Applications*, pp. 466–480, 1990.

[17] C. Slaughter, A. Y. Yang, J. Bagwell, C. Checkles, L. Sentis, and S. Vishwanath, "Sparse Online Low-rank projection and Outlier rejection (SOLO) for 3-D rigid-body motion registration." in *ICRA*, 2012.

[18] G. Sibley, C. Mei, I. Reid, and P. Newman, "Vast-scale outdoor navigation using adaptive relative bundle adjustment," *International J. of Robotics Research*, 2010.

[19] B. Kitt, A. Geiger, and H. Lategahn, "Visual odometry based on stereo image sequences with RANSAC-based outlier rejection scheme," in *IV Symposium*, 2010.

[20] D. Nistér, O. Naroditsky, and J. Bergen, "Visual odometry," in *CVPR*, vol. 1, 2004.

[21] H. Strasdat, A. J. Davison, J. M. M. Montiel, and K. Konolige, "Double window optimisation for constant time visual SLAM," in *ICCV*, 2011.

[22] C. Tomasi and T. Kanade, "Shape and motion from image streams under orthography: A factorization method," *Int. J. Comput. Vision*, 1992.

[23] Z. Lin, M. Chen, L. Wu, and Y. Ma, "The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices," UIUC, Tech. Rep., 2009.

[24] T. Bouwmans and E. Zahzah, "Robust PCA via Principal Component Pursuit: A review for a comparative evaluation in video surveillance," *Computer Vision and Image Understanding (special issue)*, 2014.

[25] Z. Lin, R. Liu, and Z. Su, "Linearized alternating direction method with adaptive penalty for low-rank representation," in *NIPS*, 2011.

[26] Y. Nesterov, "Smooth minimization of non-smooth functions," *Math. Program.*, 2005.

[27] E. Hale, W. Yin, and Y. Zhang, "Fixed-point continuation for ℓ_1-minimization: Methodology and convergence," *SIAM Journal on Optimization*, 2008.

[28] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2003.

[29] G. Ros, J. Guerrero, A. D. Sappa, D. Ponsa, and A. M. López, "VSLAM pose initialization via Lie groups and Lie algebras optimization," in *Proceedings of the IEEE ICRA*, Karlsruhe, Germany, 2013.

German Ros received his B.Sc (Hons.) in computer science from Universidad de Murcia (Spain) in 2010 and his M.Sc degree from Kingston University of London (UK) in 2011. In 2012 he obtained a second M.Sc degree from Universitat Autònoma de Barcelona, where he is currently pursuing a Ph.D. His research interests span computer vision, robotics and visual perception, with special interest in manifold optimization, robust estimation and visual geometry. He is a student member of the IEEE.

Jose M. Álvarez is currently a senior researcher at NICTA and a research fellow at the Australian National University. Previously, he was a postdoctoral researcher at the Computational and Biological Learning Group at New York University. During his Ph.D. he was a visiting researcher at the University of Amsterdam and VolksWagen research. His main research interests include road detection, color, photometric invariance, machine learning, and fusion of classifiers. He is associate editor for the IEEE Trans. on Intelligent Transportation Systems and member of the IEEE.

Julio Guerrero is an Associate Professor in the Department of Applied Maths at Universidad de Murcia (Spain) since 2000. Previously he was a postdoctoral researcher at Naples University, and visiting researcher at Durhan University (UK) and Syracure Univerity (USA). His research interests include, among others, applied maths and computer vision, with special interest in manifold optimization and robust estimation.